

Sir:

Transmitted herewith for filing is the Patent Application of:

Inventor(s): Charles Moore

For: PROCESSOR AND METHOD OF EXECUTING LOAD INSTRUCTIONS OUT-OF-ORDER HAVING REDUCED HAZARD PENALTY

Enclosed are:

☒ Patent Specification and Declaration

☒ 5 sheets of drawing(s).

☒ An assignment of the invention to International Business Machines Corporation (includes Recordation Form Cover Sheet).

☐ A certified copy of a application.

☐ Information Disclosure Statement, PTO 1449 and copies of references.

The filing fee has been calculated as shown below:

For	Number Filed	Number Extra	Rate	Fee
Basic Fee				\$690.00
Total Claims	23 - 20	3	x 18 =	\$ 54.00
Indep. Claims	5 - 3	2	x 78 =	\$156.00
MULTIPLE DEPENDENT CLAIM PRESENTED			x 260 =	\$
TOTAL				\$900.00

☒ Please charge IBM Corporation Deposit Account No. 09-0447 in the amount of \$900.00. A duplicate copy of this sheet is enclosed.

☒ The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to IBM Corporation Deposit Account 09-0447.

☒ Any additional filing fees required under 37 CFR §1.16.

☒ Any patent application processing fees under 37 CFR §1.17.

Respectfully submitted,

By Brian F. Russell

Brian F. Russell

Registration No. 40,796

FELSMAN, BRADLEY, VADEN, GUNTER & DILLON, LLP

Suite 350 Lakewood on the Park

7600B North Capital of Texas Highway

Austin, Texas 78731

Telephone (512) 343-6116

06/22/00
JC490 U.S. PTO

09/598434
06/22/00

CS11 U.S. PTO
09/598434
06/22/00

PROCESSOR AND METHOD OF EXECUTING LOAD INSTRUCTIONS OUT-
OF-ORDER HAVING REDUCED HAZARD PENALTY

BACKGROUND OF THE INVENTION

1. Technical Field:

The present invention relates in general to data processing and, in particular, to the execution of load instructions by a processor. Still more particularly, the present invention relates to a processor that buffers load data for out-of-order load instructions in order to reduce the performance penalty associated with data hazards.

2. Description of the Related Art:

A typical superscalar processor can comprise, for example, an instruction cache for storing instructions, one or more execution units for executing sequential instructions, a branch unit for executing branch instructions, instruction sequencing logic for routing instructions to the various execution units, and registers for storing operands and result data. In order to leverage the parallel execution capabilities of these multiple execution units, some superscalar processors support out-of-order execution, that is, the execution of instructions in a different order than the programmed sequence.

When executing instructions out-of-order, it is essential for correctness that the processor produce the same execution results that would have been produced had the instructions been executed in the programmed

sequence. For example, given the following sequence of instructions:

LOAD1
ADD
STORE
...
LOAD2

where LOAD1 and LOAD2 target the same address and LOAD1 precedes LOAD2 in program order, LOAD2 cannot be permitted to receive older data than LOAD1. However, if LOAD2 is executed prior to (i.e., out-of-order with respect to) LOAD1, LOAD2 may receive older data than LOAD1 if the intervening STORE is targeted at the same address or if another processor within the same computer system stores to the same address. A scenario in which an out-of-order executed load instruction receives incorrect data is defined herein to be a *data hazard*.

Superscalar processors that support out-of-order execution of load instructions typically detect and correct for data hazards by implementing a load queue that stores the target address of each load instruction that was executed out-of-order. Following execution of the out-of-order load instruction, addresses of exclusive transactions (e.g., read-with-intent-to-modify or kill) driven on the computer system interconnect by other processors, as well as store instructions preceding the load instruction that are initiated by the processor itself, are snooped against the entries within the load queue. If a snooped exclusive transaction or a local store operation hits within the load queue, the entry is marked, for example, by setting a flag.

Thereafter, when the processor executes a load instruction, the processor determines whether or not the

load instruction precedes the out-of-order load instruction in program order and whether or not the subsequently executed load instruction targets an address specified in a marked entry in the load queue. If so, a data hazard is detected, and the processor flushes and re-executes at least both load instructions, and possibly all instructions in flight following the first of the two load instruction in program order. Flushing and re-executing instructions in this manner to remedy data hazards results in a significant performance penalty, particularly for processors having wide instruction execution windows.

CONFIDENTIAL

SUMMARY OF THE INVENTION

5 The present invention reduces the performance penalty associated with data hazards resulting from the out-of-order execution of load instructions by implementing an improved load queue within a processor.

10 In accordance with the present invention, a processor having a reduced data hazard penalty includes a register set, at least one execution unit that executes load instructions to transfer data into the register set, and a load queue. The load queue contains at least one entry, and each occupied entry in the load queue stores load data retrieved by an executed load instruction in association with a target address of the executed load instruction. The load queue has associated queue management logic that, in response to execution by the execution unit of a load instruction, determines by reference to the load queue whether a data hazard exists for the load instruction. If so, the queue management logic outputs load data from the load queue to the register set in accordance with the load instruction, thus eliminating the need to flush and re-execute the load instruction.

25 All objects, features, and advantages of the present invention will become apparent in the following detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

5 The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

10 **Figure 1** depicts an illustrative embodiment of a data processing system with which the method and system of the present invention may advantageously be utilized;

15 **Figure 2** is a block diagram of an exemplary embodiment of the load data queue (LDQ) illustrated in **Figure 1**;

20 **Figure 3A** is a high level logical flowchart of an exemplary method by which the queue management logic shown in **Figure 2** updates the LDQ in response to various stages in the processing of local load operations;

25 **Figure 3B** is a high level logical flowchart of an exemplary method by which the queue management logic of **Figure 2** manages the LDQ in response to detection of local store operations and remote exclusive operations; and

30 **Figures 4A-4C** are three views of LDQ 114 that together illustrate an exemplary operating scenario in which a data hazard caused by a remote exclusive operation is detected and corrected in accordance with

the present invention.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENT

With reference now to the figures and in particular with reference to **Figure 1**, there is depicted a high level block diagram of an illustrative embodiment of a processor, indicated generally at **10**, for processing instructions and data in accordance with the present invention. In particular, processor **10** reduces data hazard penalties by implementing a load data queue that buffers load data associated with out-of-order load instructions.

Processor **10** comprises a single integrated circuit superscalar processor, which, as discussed further below, includes various execution units, registers, buffers, memories, and other functional units that are all formed by integrated circuitry. As illustrated in **Figure 1**, processor **10** may be coupled to other devices, such as a system memory **12** and a second processor **10**, by an interconnect fabric **14** to form a larger data processing system such as a workstation computer system. Processor **10** also includes an on-chip multi-level cache hierarchy including a unified level two (L2) cache **16** and bifurcated level one (L1) instruction (I) and data (D) caches **18** and **20**, respectively. As is well-known to those skilled in the art, caches **16**, **18** and **20** provide low latency access to cache lines corresponding to memory locations in system memory **12**.

Instructions are fetched for processing from L1 I-cache **18** in response to the effective address (EA) residing in instruction fetch address register (IFAR) **30**.

During each cycle, a new instruction fetch address may be loaded into IFAR 30 from one of three sources: branch prediction unit (BPU) 36, which provides speculative target path addresses resulting from the prediction of conditional branch instructions, global completion table (GCT) 38, which provides sequential path addresses, and branch execution unit (BEU) 92, which provides non-speculative addresses resulting from the resolution of predicted conditional branch instructions. If hit/miss logic 22 determines, after translation of the EA contained in IFAR 30 by effective-to-real address translation (ERAT) 32 and lookup of the real address (RA) in I-cache directory 34, that the cache line of instructions corresponding to the EA in IFAR 30 does not reside in L1 I-cache 18, then hit/miss logic 22 provides the RA to L2 cache 16 as a request address via I-cache request bus 24. Such request addresses may also be generated by prefetch logic within L2 cache 16 based upon recent access patterns. In response to a request address, L2 cache 16 outputs a cache line of instructions, which are loaded into prefetch buffer (PB) 28 and L1 I-cache 18 via I-cache reload bus 26, possibly after passing through optional predecode logic 144.

Once the cache line specified by the EA in IFAR 30 resides in L1 cache 18, L1 I-cache 18 outputs the cache line to both branch prediction unit (BPU) 36 and to instruction fetch buffer (IFB) 40. BPU 36 scans the cache line of instructions for branch instructions and predicts the outcome of conditional branch instructions, if any. Following a branch prediction, BPU 36 furnishes

a speculative instruction fetch address to IFAR 30, as discussed above, and passes the prediction to branch instruction queue 64 so that the accuracy of the prediction can be determined when the conditional branch instruction is subsequently resolved by branch execution unit 92.

IFB 40 temporarily buffers the cache line of instructions received from L1 I-cache 18 until the cache line of instructions can be translated by instruction translation unit (ITU) 42. In the illustrated embodiment of processor 10, ITU 42 translates instructions from user instruction set architecture (UISA) instructions into a possibly different number of internal ISA (IISA) instructions that are directly executable by the execution units of processor 10. Such translation may be performed, for example, by reference to microcode stored in a read-only memory (ROM) template. In at least some embodiments, the UISA-to-IISA translation results in a different number of IISA instructions than UISA instructions and/or IISA instructions of different lengths than corresponding UISA instructions. The resultant IISA instructions are then assigned by global completion table 38 to an instruction group, the members of which are permitted to be executed out-of-order with respect to one another. Global completion table 38 tracks each instruction group for which execution has yet to be completed by at least one associated EA, which is preferably the EA of the oldest instruction in the instruction group.

Following UISA-to-IISA instruction translation, instructions are dispatched in-order to one of latches

44, 46, 48 and 50 according to instruction type. That is, branch instructions and other condition register (CR) modifying instructions are dispatched to latch 44, fixed-point and load-store instructions are dispatched to either of latches 46 and 48, and floating-point instructions are dispatched to latch 50. Each instruction requiring a rename register for temporarily storing execution results is then assigned one or more rename registers by the appropriate one of CR mapper 52, link and count (LC) register mapper 54, exception register (XER) mapper 56, general-purpose register (GPR) mapper 58, and floating-point register (FPR) mapper 60.

The dispatched instructions are then temporarily placed in an appropriate one of CR issue queue (CRIQ) 62, branch issue queue (BIQ) 64, fixed-point issue queues (FXIQs) 66 and 68, and floating-point issue queues (FPIQs) 70 and 72. From issue queues 62, 64, 66, 68, 70 and 72, instructions can be issued opportunistically (i.e., possibly out-of-order) to the execution units of processor 10 for execution. The instructions, however, are maintained in issue queues 62-72 until execution of the instructions is complete and the result data, if any, are written back, in case any of the instructions needs to be reissued.

As illustrated, the execution units of processor 10 include a CR unit (CRU) 90 for executing CR-modifying instructions, a branch execution unit (BEU) 92 for executing branch instructions, two fixed-point units (FXUs) 94 and 100 for executing fixed-point instructions,

two load-store units (LSUs) 96 and 98 for executing load and store instructions, and two floating-point units (FPUs) 102 and 104 for executing floating-point instructions. Each of execution units 90-104 is preferably implemented as an execution pipeline having a number of pipeline stages.

During execution within one of execution units 90-104, an instruction receives operands, if any, from one or more architected and/or rename registers within a register file coupled to the execution unit. When executing CR-modifying or CR-dependent instructions, CRU 90 and BEU 92 access the CR register file 80, which in a preferred embodiment contains a CR and a number of CR rename registers that each comprise a number of distinct fields formed of one or more bits. Among these fields are LT, GT, and EQ fields that respectively indicate if a value (typically the result or operand of an instruction) is less than zero, greater than zero, or equal to zero. Link and count register (LCR) register file 82 contains a count register (CTR), a link register (LR) and rename registers of each, by which BEU 92 may also resolve conditional branches to obtain a path address. General-purpose register files (GPRs) 84 and 86, which are synchronized, duplicate register files, store fixed-point and integer values accessed and produced by FXUs 94 and 100 and LSUs 96 and 98. Floating-point register file (FPR) 88, which like GPRs 84 and 86 may also be implemented as duplicate sets of synchronized registers, contains floating-point values that result from the execution of floating-point instructions by FPUs 102 and 104 and floating-point load instructions by LSUs 96 and

98.

After an execution unit finishes execution of an instruction, the execution notifies GCT 38, which schedules completion of instructions in program order. To complete an instruction executed by one of CRU 90, FXUs 94 and 100 or FPUs 102 and 104, GCT 38 signals the execution unit, which writes back the result data, if any, from the assigned rename register(s) to one or more architected registers within the appropriate register file. The instruction is then removed from the issue queue, and once all instructions within its instruction group have completed, is removed from GCT 38. Other types of instructions, however, are completed differently.

When BEU 92 resolves a conditional branch instruction and determines the path address of the execution path that should be taken, the path address is compared against the speculative path address predicted by BPU 36. If the path addresses match, no further processing is required. If, however, the calculated path address does not match the predicted path address, BEU 92 supplies the correct path address to IFAR 30. In either event, the branch instruction can then be removed from BIQ 64, and when all other instructions within the same instruction group have completed, from GCT 38.

Following execution of a load instruction, the effective address computed by executing the load instruction is translated to a real address by a data ERAT (not illustrated) and then provided to L1 D-cache

as a request address. At this point, the load operation is removed from FXIQ 66 or 68 and placed in load data queue (LDQ) 114 until the indicated load is performed. If the request address misses in L1 D-cache 20, the request address is placed in load miss queue (LMQ) 116, from which the requested data is retrieved from L2 cache 16, and failing that, from another processor 10 or from system memory 12. As described in detail below, LDQ 114 ensures that data hazards are detected and appropriate remedial action is taken such that the later of two load instructions targeting the same address does not receive older data than the earlier of the two load instructions. Store instructions are similarly completed utilizing a store queue (STQ) 110 into which effective addresses for stores are loaded following execution of the store instructions. From STQ 110, data can be stored into either or both of L1 D-cache 20 and L2 cache 16.

Referring now to **Figure 2**, there is depicted an exemplary embodiment of LDQ 114 of processor 10. As illustrated, LDQ 114 includes a number of entries, each including a effective address (EA) field 120 for storing the effective address (of address tag portion thereof) of a load instruction, a target address field 122 for storing the target address (or address tag portion thereof) from which the load instruction obtains data, a data field 124 for storing data loaded from memory by a load instruction, and a hazard field 126 for indicating that a hazard may exist for a load instruction. Entries within LDQ 114 are preferably allocated, updated, and deallocated by associated queue management logic 128 in

accordance with the processes depicted in **Figures 3A** and **3B**.

With reference now to **Figure 3A**, there is illustrated a high level logical flowchart of an exemplary method by which queue management logic **128** of **Figure 2** manages LDQ **114** in response to various stages in the local processing of load operations. As shown, the process begins at block **130** and then proceeds to block **132** in response to queue management logic **128** receiving a notification that a load instruction has been processed at some stage of the execution pipeline between dispatch and completion. In response to this notification, queue management logic **128** determines at block **132** whether the load instruction has been dispatched, executed or completed by processor **10**. In response to a determination that the load instruction has been dispatched from ITU **42** to one of latches **46** and **48**, as described above, the process proceeds to block **134**. Block **134** depicts queue management logic **126** allocating an entry in LDQ **114** for the newly dispatched instruction in accordance with the program order of the load instruction and placing the EA of the instruction within EA field **120**. Thus, the location of an entry of a load instruction within LDQ **114** preferably indicates the program ordering of the load instruction with respect to other load instructions. Thereafter, the process returns to block **132**.

Returning to block **132**, in response to a determination that a load instruction has been completed

(together with other instructions in its instruction group) by GCT 38, queue management logic 126 deallocates the entry corresponding to the completed load instruction, for example, by identifying an entry having a matching EA. Thereafter, the process returns to block 132.

If, on the other hand, queue management logic 128 determines from the received notification at block 132 that a load instruction has been executed by one of LSUs 96 and 98, the process proceeds to block 140, which illustrates queue management logic 128 determining whether a later entry in LDQ 114 than the entry allocated to the executed load instruction has a target address in its target address field 122 that matches the target address of the executed load instruction.

If not, queue management logic 128 places the target address of the executed load instruction in the target address field 122 of the associated entry and places the data retrieved from memory (i.e., local cache, remote cache, or system memory 12) in response to execution of the load instruction in data field 124 of the associated entry, as shown at block 142. The entry associated with the executed load instruction is also updated, as depicted at block 142, even if an entry associated with a later load instruction has a matching address if a determination is made at block 144 that hazard field 126 of the matching entry is not set. However, if hazard field 126 of the matching entry is set, a data hazard is detected.

As illustrated at block 146, to correct for the data hazard, queue management logic 128 places the target address for the executed load instruction in target address field 122 of the associated entry and utilizes the data contained in data field 124 of the matching entry of the later-in-program-order load to provide the data requested by the executed load instruction. That is, the data from data field 124 of the matching entry is provided to one of GPRs 84 and 86 as specified by the executed load instruction and is also placed into data field 124 of the entry in LDQ 114 associated with the executed load instruction. Thus, the operation of queue management logic 128 minimizes the performance penalty associated with data hazards since the earlier-in-program-order load instruction need not be re-executed to obtain the correct data (i.e., in this case, the same data as the later-in-program-order load) and no flush of instructions is required. Following block 146, the process returns to block 132.

Referring now to **Figure 3B**, there is depicted a high level logical flowchart of an exemplary method by which queue management logic 128 of **Figure 2** updates hazard fields 126 of LDQ 114 in response to detection of remote exclusive operations and corrects data hazards occasioned by the execution of local store operations. As depicted, the process begins at block 150 and then iterates at block 152 until notification is received of a locally executed store instruction or of an exclusive access (e.g., read-with-intent-to-modify, flush or kill) request on interconnect fabric 14 made by a remote

processor 10. In response to receipt of notification of local execution of a store instruction, the process passes to block 160, which is described below. However, in response to notification of an exclusive access request by a remote processor 10, the process proceeds from block 152 to block 154, which illustrates queue management logic 128 determining whether or not a target address specified by the remote exclusive access request matches the target address contained in target address field 122 of any entry within LDQ 114. If not, the process simply returns to block 152, which has been described.

However, in response to a determination that the target address of the remote exclusive address request matches the address contained in the target address field 122 of an entry in LDQ 114, queue management logic 128 sets hazard field 126 of the matching entry, as shown at block 156, to indicate the existence of a possible remotely-triggered data hazard for any earlier-in-program-order load instruction executed after the load instruction associated with matching entry. The existence of an actual data hazard is detected at blocks 140 and 144 of **Figure 3A** for remote exclusive operations and at blocks 160-164 of **Figure 3B** for local store operations. Following block 156, the process illustrated in **Figure 3B** returns to block 152.

Referring again to block 152, in response to queue management logic 128 receiving notification of execution of a local store instruction, queue management

logic 128 determines at blocks 160 and 162 whether or not the target address of the store instruction matches a target address of an later-in-program-order but earlier executed load instruction in one of target address fields 122 of LDQ 114. If not, the process simply returns to block 152, which has been described. However, in response to a determination that the target address of the store instruction matches a target address of an later-in-program-order but earlier executed load instruction, queue management logic 128 determines that a data hazard has occurred and corrects the data hazard by flushing at least the matching load instruction and any subsequent dependent instructions and by causing these instructions to be re-executed. Queue management logic 128 also deallocates the entry in LDQ 114 allocated to the flushed load.

With reference now to **Figures 4A-4C**, there are illustrated three block diagrams that together illustrate an exemplary operating scenario in which a data hazard caused by a remote exclusive operation is detected and corrected in accordance with the present invention. Referring first to **Figure 4A**, when the operating scenario begins, two load instructions, which are designated LD1 and LD2 in program order (with LD2 being the latest in program order), have been dispatched and accordingly have been allocated entries in LDQ 114 by queue management logic 128. In addition, LD2 has been executed out-of-order with respect to LD1, and the target address (TA) and data (D) have been loaded into the appropriate entry of LDQ 114 by queue management logic 128. The hazard field 126 of the entry associated with each of the load

instructions is reset to 0.

5 Next, as shown in **Figure 4B**, in response to
queue management logic 128 receiving notification of a
remote request for exclusive access having a target
address that matches the TA of LD2, hazard field 126 of
the entry associated with LD2 is set to 1. Then, as
indicated in **Figure 4C**, when LD1 is executed out-of-order
and the execution generates a target address matching the
10 TA specified in target address field 124 of the entry
associated with LD2, a data hazard is detected.
Accordingly, queue management logic 128 provides the data
from data field 124 of the entry corresponding to LD2 to
the register file to satisfy LD1 and also records the
15 data in data field 124 and records the target address in
target address field 122 of the entry corresponding to
LD1. Thus, a data hazard caused by a remote exclusive
operation intervening between out-of-order executed loads
is detected and correctly without flushing or re-
20 executing any instructions and without any additional
latency.

25 As has been described, the present invention
provides an improved processor and method that reduces
the performance penalty associated with data hazards by
recording the data associated with out-of-order load
instructions in a load data queue and then, in response
to detection of a data hazard, utilizing that data to
satisfy an earlier-in-program-order load instruction.

30 While the invention has been particularly shown
and described with reference to a preferred embodiment,
it will be understood by those skilled in the art that

various changes in form and detail may be made therein without departing from the spirit and scope of the invention.

20230714 14:00:00

CLAIMS

What is claimed is:

1. A method of processing an instruction in a processor,
said method comprising:

executing a first instruction and storing data
associated with executing said first instruction;

if said first instruction was executed out-of-
order, determining whether the stored data should have
been used in execution of a preceding second instruction
executed after said first instruction; and

if the stored data should have been used to
execute the second instruction, using the stored data to
perform an operation indicated by said second instruction
without re-executing said second instruction.

1 2. A processor, comprising:

2 means for executing a first instruction and
3 storing data associated with executing said first
4 instruction;

5 means for determining, if said first
6 instruction was executed out-of-order, whether the stored
7 data should have been used in execution of a preceding
8 second instruction executed after said first instruction;
9 and

10 means for using the stored data to perform an
11 operation indicated by said second instruction without
12 re-executing said second instruction, if the stored data
13 should have been used to execute the second instruction.

1 3. A processor, comprising:

2 a register set;

3 at least one execution unit that executes load
4 instructions to transfer data into said register set;

5 a load queue containing at least one entry,
6 wherein said entry stores load data retrieved by a first
7 load instruction; and

8 queue management logic that, responsive to
9 execution of a second load instruction, detects by
10 reference to said load queue whether a data hazard
11 exists, and if so, outputs said load data retrieved by
12 said first load instruction from said entry to said
13 register set in accordance with said second load
14 instruction.

1 4. The processor of Claim 3, wherein said entry stores
2 a target address of said first load instruction and has a
3 hazard flag indicative of a possible data hazard, wherein
4 said queue management logic detects that a data hazard
5 exists if said second load instruction precedes said
6 first instruction in program order and a target address
7 of said second load instruction matches said target
8 address stored in said entry and said hazard flag is set.

1 5. The processor of Claim 3, wherein said queue
2 management logic sets said hazard flag at least in
3 response to local store operation specifying said target
4 address.

1 6. The processor of Claim 3, said register set
2 comprising a general purpose register set.

1 7. The processor of Claim 3, wherein said queue
2 management logic outputs said load data to a register in
3 said register set that is specified by said second load
4 instruction.

1 8. The processor of Claim 3, wherein said queue
2 management logic, responsive to detection of a data
3 hazard, initiates reexecution of at least said first load
4 instruction.

1 9. The processor of Claim 3, wherein said queue
2 management logic allocates a respective entry within said
3 load queue to each load instruction upon dispatch and,
4 upon completion of said each load instruction,
5 deallocates said respective entry.

1 10. A data processing system, comprising:

2 an interconnect fabric;

3 a memory coupled to said interconnect fabric;

4 a register set;

5 at least one execution unit that executes load
6 instructions to transfer data from said memory into said
7 register set;

8 a load queue containing at least one entry,
9 wherein said entry stores load data retrieved by a first
10 load instruction; and

11 queue management logic that, responsive to
12 execution of a second load instruction, detects by
13 reference to said load queue whether a data hazard
14 exists, and if so, outputs said load data retrieved by
15 said first load instruction from said entry to said
16 register set in accordance with said second load
17 instruction.

1 11. The data processing system of Claim 10, wherein said
2 entry stores a target address of said first load
3 instruction and has a hazard flag indicative of a
4 possible data hazard, wherein said queue management logic
5 detects that a data hazard exists if said second load
6 instruction precedes said first instruction in program
7 order and a target address of said second load
8 instruction matches said target address stored in said
9 entry and said hazard flag is set.

1 12. The data processing system of Claim 11, wherein said
2 queue management logic sets said hazard flag at least in
3 response to local store operation specifying said target
4 address.

1 13. The data processing system of Claim 12, wherein said
2 at least one execution unit, said register set and said
3 load queue comprise a first processor and said data
4 processing system includes a second processor, wherein
5 said queue management logic also sets said hazard flag in
6 response to said second processor issuing an exclusive
7 access operation specifying said target address on said
8 interconnect fabric.

1 14. The data processing system of Claim 10, said
2 register set comprising a general purpose register set.

1 15. The data processing system of Claim 10, wherein said
2 queue management logic outputs said load data to a
3 register in said register set that is specified by said
4 second load instruction.

1 16. The data processing system of Claim 10, wherein said
2 queue management logic, responsive to detection of a data
3 hazard, initiates reexecution of at least said first load
4 instruction.

1 17. The data processing system of Claim 10, wherein said
2 queue management logic allocates a respective entry
3 within said load queue to each load instruction upon
4 dispatch and, upon completion of said each load
5 instruction, deallocates said respective entry.

1 18. A method of executing load instructions out-of-order
2 in a processor having a register set and a load queue,
3 said method comprising:

4 storing, in an entry in said load queue, load
5 data retrieved from memory in response to executing a
6 first load instruction;

7 in response to execution of a second load
8 instruction, detecting by reference to said load queue
9 whether a data hazard exists; and

10 in response to detection of a data hazard,
11 outputting said load data retrieved by said first load
12 instruction from said entry to said register set in
13 accordance with said second load instruction.

1 19. The method of Claim 18, wherein said entry stores a
2 target address of said first load instruction and has a
3 hazard flag indicative of a possible data hazard, wherein
4 detecting that a data hazard exists comprises determining
5 if said second load instruction precedes said first
6 instruction in program order and a target address of said
7 second load instruction matches said target address
8 stored in said entry and said hazard flag is set.

1 20. The method of Claim 19, and further comprising
2 setting said hazard flag at least in response to a local
3 store operation specifying said target address.

1 21. The method of Claim 19, wherein outputting said load
2 data comprises outputting said load data to a register in
3 said register set that is specified by said second load
4 instruction.

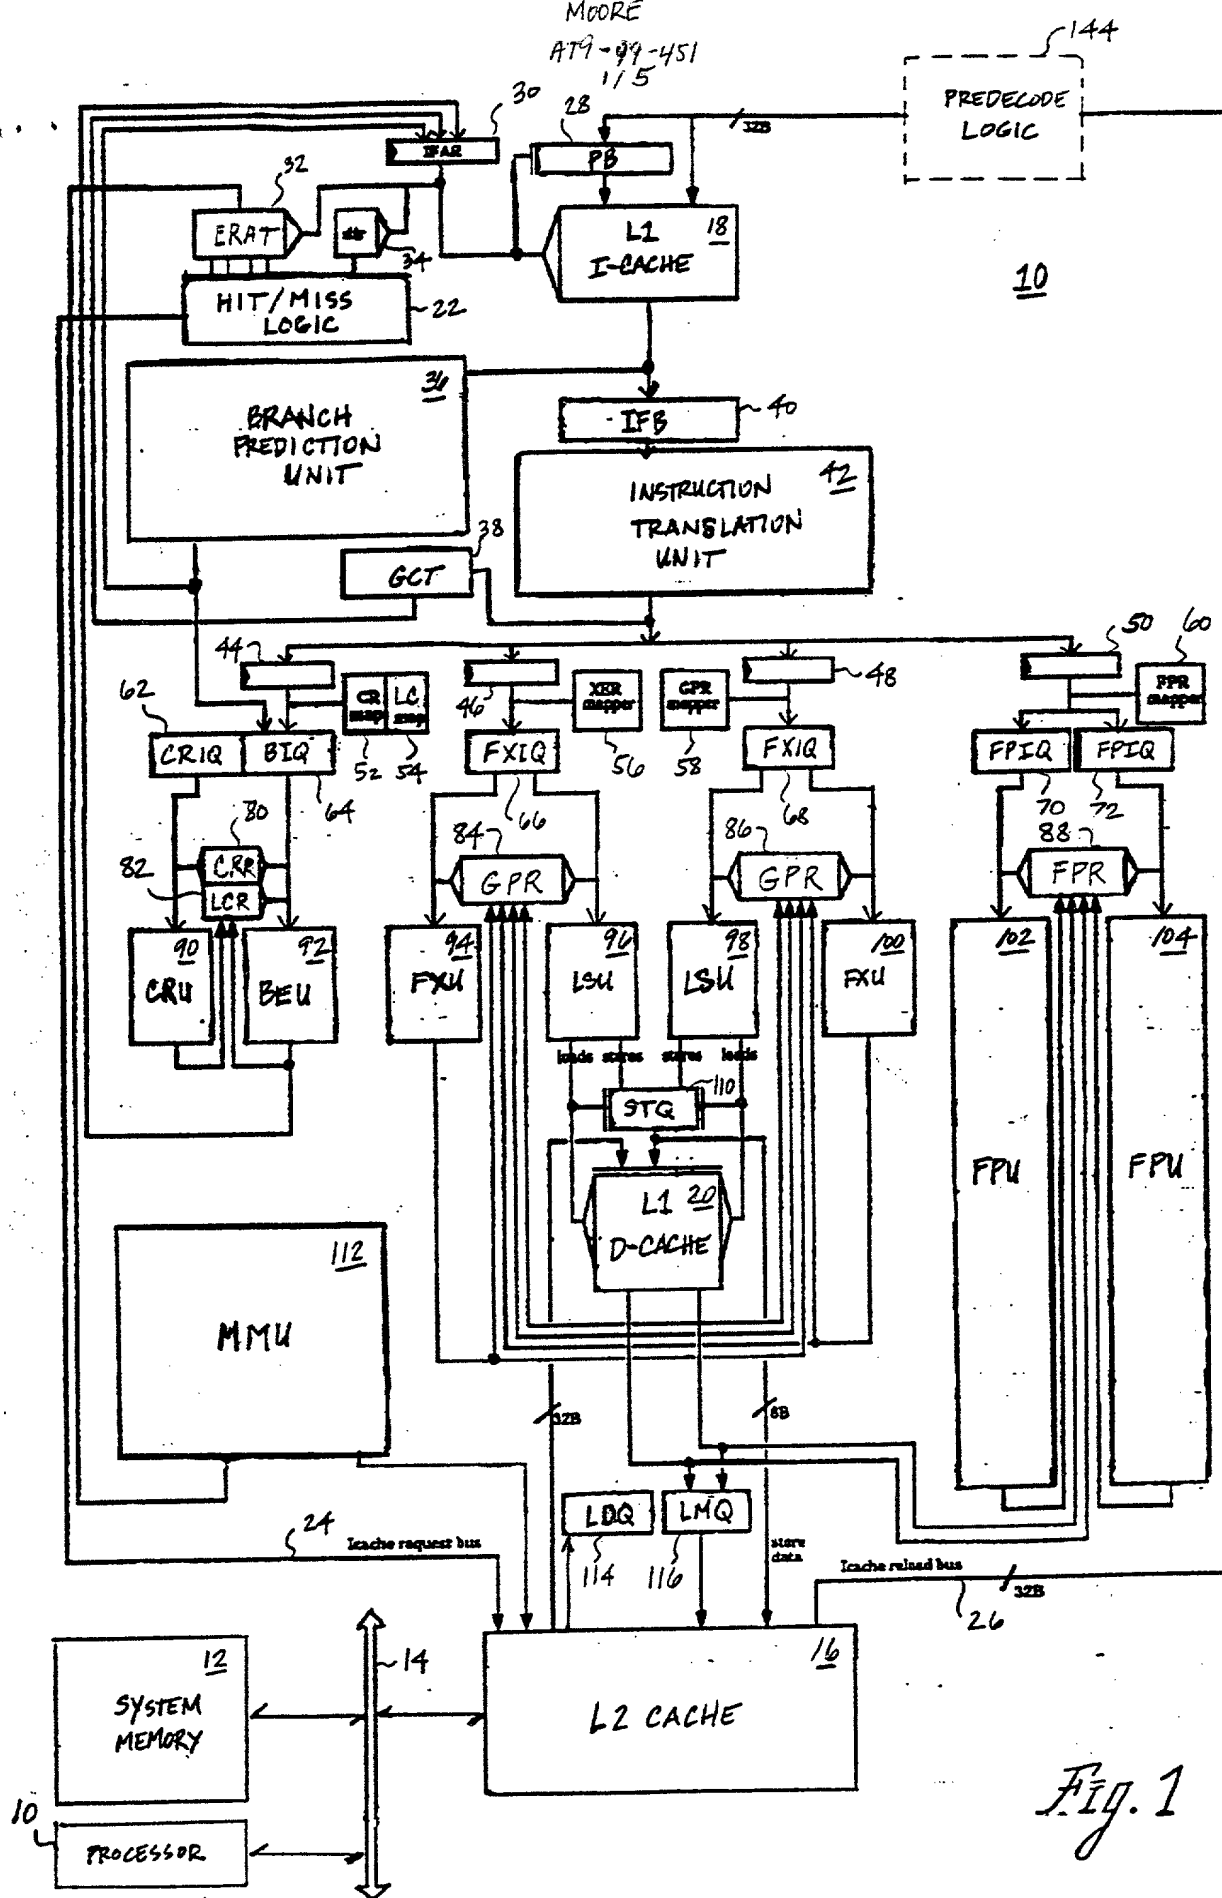
1 22. The method of Claim 19, and further comprising:

2 in response to detection of a data hazard,
3 initiating reexecution of at least said first load
4 instruction.

1 23. The method of Claim 19, and further comprising
2 allocating a respective entry within said load queue to
3 each load instruction upon dispatch and, upon completion
4 of said each load instruction, deallocating said
5 respective entry.

ABSTRACT OF THE DISCLOSURE
PROCESSOR AND METHOD OF EXECUTING LOAD INSTRUCTIONS OUT-
OF-ORDER HAVING REDUCED HAZARD PENALTY

5 A processor having a reduced data hazard
penalty includes a register set, at least one execution
unit that executes load instructions to transfer data
into the register set, and a load queue. The load queue
contains at least one entry, and each occupied entry in
the load queue stores load data retrieved by an executed
load instruction in association with a target address of
10 the executed load instruction. The load queue has
associated queue management logic that, in response to
execution by the execution unit of a load instruction,
determines by reference to the load queue whether a data
hazard exists for the load instruction. If so, the queue
management logic outputs load data from the load queue to
15 the register set in accordance with the load instruction,
thus eliminating the need to flush and re-execute the
load instruction.






Fig. 2

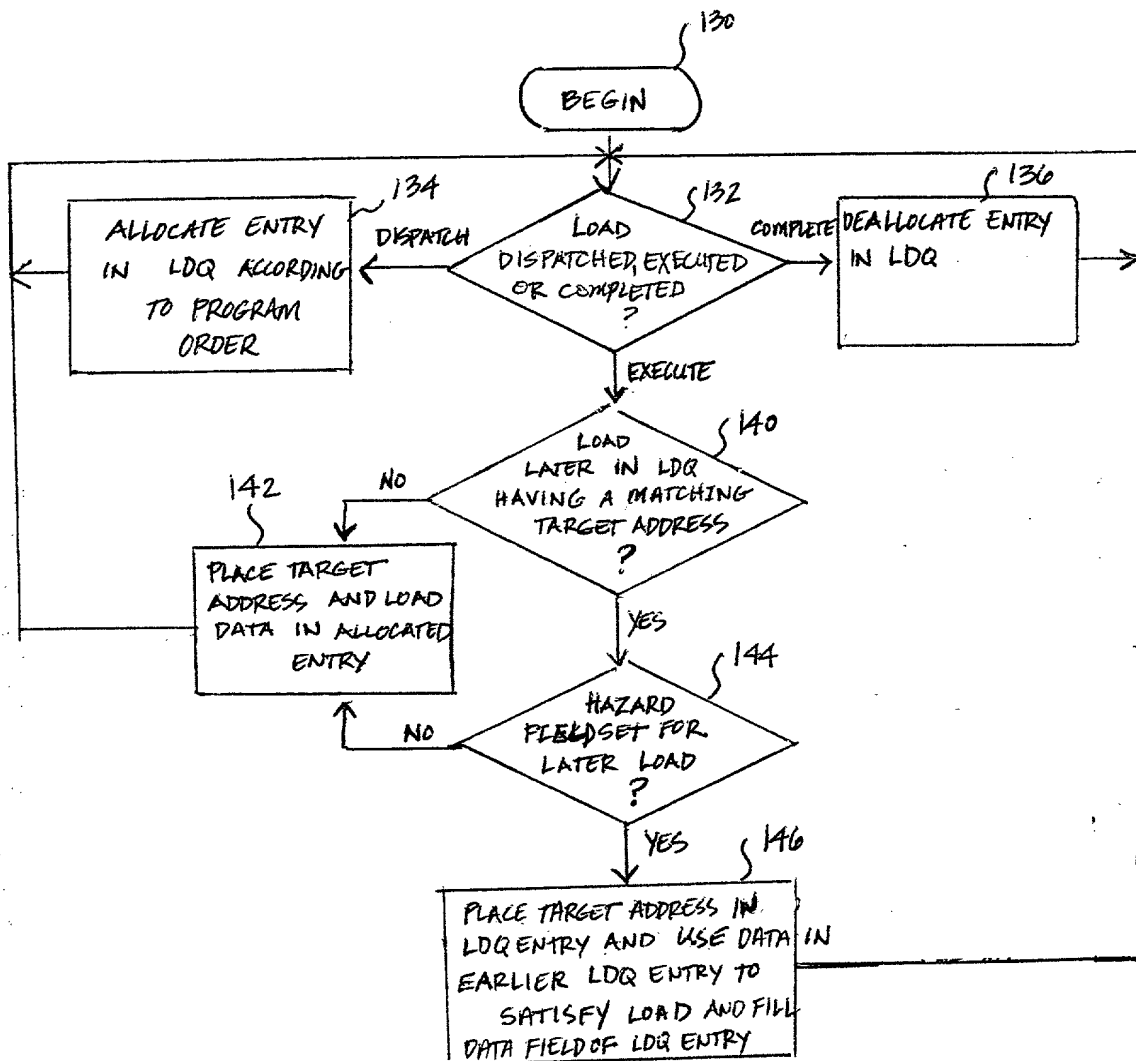


Fig. 3A

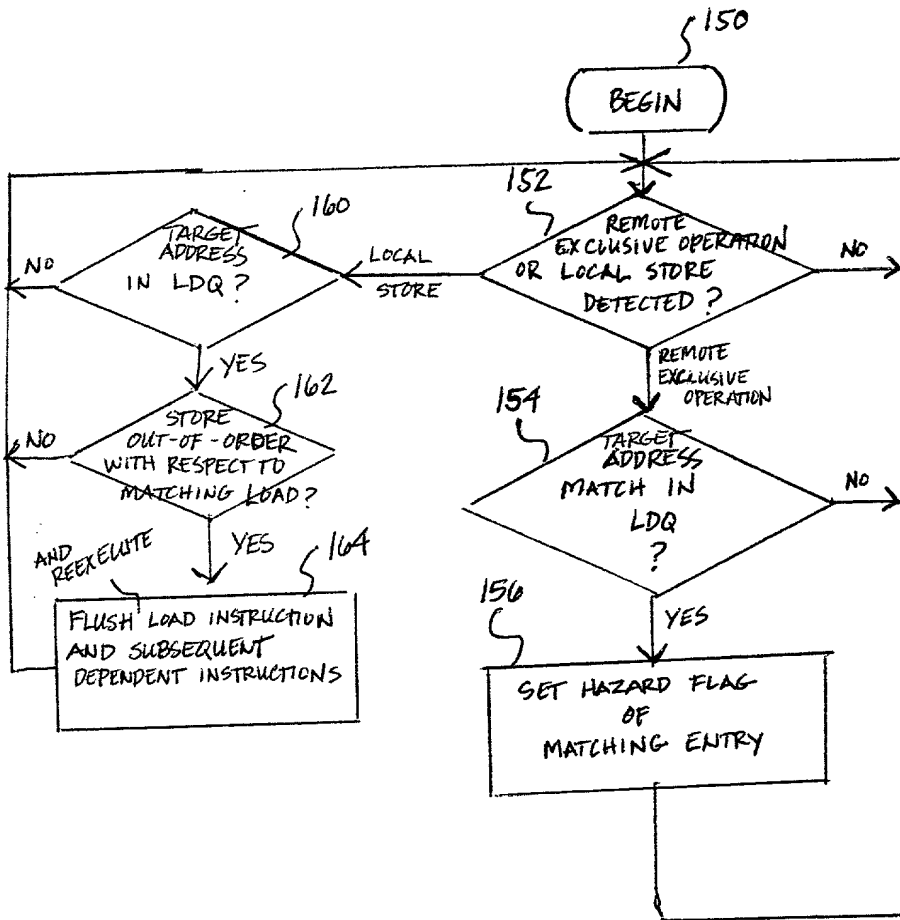


Fig. 3B




Fig. 4A




Fig. 4B




Fig. 4C

**DECLARATION AND POWER OF ATTORNEY FOR
PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

PROCESSOR AND METHOD OF EXECUTING LOAD INSTRUCTIONS OUT-OF-ORDER HAVING REDUCED HAZARD PENALTY

the specification of which (check one)

X is attached hereto.

___ was filed on _____
as Application Serial No. _____
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s):

Priority Claimed

____ Yes ____ No

(Number) (Country) (Day/Month/Year)

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, Code of Federal

Regulations, \$1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial #)	(Filing Date)	(Status)
------------------------	---------------	----------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

John W. Henderson, Jr., Reg. No. 26,907; Thomas E. Tyson, Reg. No. 28,543; Robert M. Carwell, Reg. No. 28,499; Jeffrey S. LaBaw, Reg. No. 31,633; Douglas H. Lefevre, Reg. No. 26,193; Casimer K. Salys, Reg. No. 28,900; David A. Mims, Jr., Reg. No. 32,708; Volel Emile, Reg. No. 39,969; James H. Barksdale, Jr. Reg. No. 24,091; Anthony V. England, Reg. No. 35,129; Leslie A. Van Leeuwen, Reg. No. 42,196; Marilyn S. Dawkins, Reg. No. 31,140; Mark E. McBurney, Reg. No. 33,114; Christopher A. Hughes, Reg. No. 26,914; Edward A. Pennington, Reg. No. 32,588; John E. Hoel, Reg. No. 26,279; Joseph C. Redmond, Jr., Reg. No. 18,753; Andrew J. Dillon, Reg. No. 29,634; Max Ciccarella, Reg. No. 39,454; Daniel E. Venglarik, Reg. No. 39,409; Jack V. Musgrove, Reg. No. 31,986; Brian F. Russell, Reg. No. 40,796; Steven Lin, Reg. No. 35,250; Matthew W. Baca, Reg. No. 42,277; Antony P. Ng, Reg. No. 43,427; John G. Graham, Reg. No. 19,563; Matthew S. Anderson, Reg. No. 39,093; Michael R. Barre, Reg. No. 44,023; Andrew Mitchell Harris, Reg. No. 42,638; Richard McCain, Reg. No. 43,785; Michael Noe, Reg. No. 44,975; and Sidney L. Weatherford, Reg. No. P-45,602.

Send correspondence to: Andrew J. Dillon, FELSMAN, BRADLEY, VADEN, GUNTER & DILLON, LLP, Suite 350 Lakewood on the Park, 7600B North Capital of Texas Highway, Austin, Texas 78731, and direct all telephone calls to Andrew J. Dillon, (512) 343-6116.

FULL NAME OF SOLE OR FIRST INVENTOR: Charles Robert Moore

INVENTORS SIGNATURE: Charles Robert Moore DATE: 6/12/2000

RESIDENCE: 8802 Royalwood Drive
Austin, Texas 78750

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 8802 Royalwood Drive
Austin, Texas 78750